arXiv:1612.08499v1 [cs.CV] 27 Dec 2016

End-to-End Data Visualization by Metric Learning and Coordinate Transformation

Lilei Zheng[1]
zhengll@i2r.a-star.edu.sg
Ying Zhang[2]
zhangy@i2r.a-star.edu.sg
Stefan Duffner[1]
stefan.duffner@liris.cnrs.fr
Khalid Idrissi[1]
khalid.idrissi@liris.cnrs.fr
Christophe Garcia[1]
christophe.garcia@liris.cnrs.fr
Atilla Baskurt[1]
atilla.baskurt@liris.cnrs.fr

[1] Université de Lyon,
CNRS, INSA-Lyon, LIRIS,
UMR5205, F-69621, France
[2] Cyber Security & Intelligence,
Institute for Infocomm Research,
A*STAR, Singapore

Abstract

This paper presents a deep nonlinear metric learning framework for data visualization on an image dataset. We propose the Triangular Similarity and prove its equivalence to the Cosine Similarity in measuring a data pair. Based on this novel similarity, a geometrically motivated loss function – the triangular loss – is then developed for optimizing a metric learning system comprising two identical CNNs. It is shown that this deep nonlinear system can be efficiently trained by a hybrid algorithm based on the conventional backpropagation algorithm. More interestingly, benefiting from classical manifold learning theories, the proposed system offers two different views to visualize the outputs, the second of which provides better classification results than the state-of-the-art methods in the visualizable spaces.

1 Introduction

End-to-end learning trains neural networks directly on raw collected data such as speech or image to achieve effective feature learning for classification or other tasks [17, 23]. Thus it has recently received much attention because it requires no handcrafted feature extraction algorithms but can directly predict outputs from low-level inputs. For example, for speech recognition, Recurrent Neural Networks (RNN) learn to map from acoustic signals to phonetic sequences [17]; for image classification, Convolutional Neural Networks (CNN) are originally designed to receive an image as input and can classify raw pixels into high-level concepts such as object categories [23].

Traditional supervised neural networks learn on labeled data samples and require handcrafted target vectors to conduct error backpropagation, so a CNN classifier is only available

© 2016. The copyright of this document resides with its authors.

for restricted dimensionality reduction because the size of the output layer (i.e. the target dimension) is fixed to the number of classes in a classification problem. This restriction can be relaxed by putting neural networks in a *siamese architecture* [6] which consists of two identical sub-systems sharing the same set of parameters. This technique, namely Siamese Neural Networks (SNN), performs learning from relationship between data pairs instead of directly on labeled data samples. The SNN approach is categorized to Metric Learning (ML) [3], i.e. methods that automatically learn a metric from a set of data pairs. Therefore, this work proposes a novel metric learning method for end-to-end learning on images.

A good Metric Learning system should be a collaborative product of designing a mapping architecture and formulating a metric-based cost function [3, 11]. For a mapping function, the depth of the architecture can be shallow or deep, namely, shallow linear transformations [3], multi-layer neural networks [5, 10, 36, 42] or deep neural networks [1, 6, 11, 19, 30]; for a cost function, it may be based on a distance metric [12, 14, 16, 18, 37, 40, 41], a similarity metric [9, 31, 34, 35] or a hybrid metric concerning both distance and similarity [2, 7].

As a commonly used metric for measuring the similarity of two feature vectors, the Cosine Similarity and its variant the bilinear similarity have been found particularly useful for text retrieval [8, 9, 34, 35] and face verification [7, 31]. This paper contributes to the study of similarity metric learning and proposes an alternative metric which is equivalent to the Cosine Similarity in measuring a data pair but has a nicer geometrical interpretation of learning the similarity. This novel metric is naturally related to the well-known *triangle inequality theorem*, so we call it the *Triangular Similarity*. The present definition of Triangular Similarity provides a theoretical explanation of the *triangular loss* function that was first given in [45]. Moreover, by examining the gradient of this cost function, we discover its relationship to the Mean Square Error (MSE) function of the traditional neural networks. Therefore typical optimization techniques of training neural networks can be directly applied to optimize our metric learning system.

We integrate the triangular loss function with two CNNs in a siamese architecture, to relax the constraint on the output dimension and make flexible dimensionality reduction to the input data. We name this method as CNN-based Triangular Similarity Metric Learning (TSML-CNN). The method is tested on mapping raw handwritten digit images of the MNIST dataset into visualizable spaces, i.e. target spaces of dimension no larger than 3. This particular kind of dimensionality reduction is so called *data visualization*. The contributions of this paper are the following.

1) a theoretical proof of learning the Triangular Similarity is explained and illustrated. By a soft length normalization term, we deduce the triangular loss function from the simplified Triangular Similarity. This completes the theory of TSML.

2) we present TSML-CNN to perform flexible dimensionality reduction. Compared with that of training a traditional CNN, a significant increase of training time is observed for this siamese variant. However, we introduce a hybrid training strategy to accelerate the training speed.

3) the proposed method realizes end-to-end data visualization *without* losing its superior ability of accurate classification. It achieves competitive performance on class identification in low dimensional spaces.

4) we introduce a novel perspective for data visualization in manifold learning field by simple coordinate transformation on the neural networks outputs. Thus we can view

the data either on the unit sphere (hypersphere) in a target space, or on the unfolded plane (hyperplane) in an even lower space.

2 Triangular Similarity Metric Learning

2.1 Triangular Similarity

For any two given vectors $\mathbf{a}$ and $\mathbf{b}$, the Triangular Similarity between them is measured by:

$$tri(\mathbf{a},\mathbf{b}) = \frac{1}{2}\|\frac{\mathbf{a}}{\|\mathbf{a}\|} + \frac{\mathbf{b}}{\|\mathbf{b}\|}\|, \tag{1}$$

the value of this similarity lies in the range [0, 1]: it results in 1 if and only if the two vectors $\mathbf{a}$ and $\mathbf{b}$ are towards the same direction, and yields 0 if and only if the directions of the two vectors are exactly opposite. The Triangular Similarity is related to the Cosine Similarity by:

$$\begin{aligned} tri(\mathbf{a},\mathbf{b}) &= \frac{1}{2}\|\frac{\mathbf{a}}{\|\mathbf{a}\|} + \frac{\mathbf{b}}{\|\mathbf{b}\|}\| = \frac{1}{2}\sqrt{(\frac{\mathbf{a}}{\|\mathbf{a}\|})^2 + (\frac{\mathbf{b}}{\|\mathbf{b}\|})^2 + 2\frac{\mathbf{a}^T\mathbf{b}}{\|\mathbf{a}\|\|\mathbf{b}\|}} \\ &= \frac{1}{2}\sqrt{2+2cos(\mathbf{a},\mathbf{b})} = \sqrt{\frac{1+cos(\mathbf{a},\mathbf{b})}{2}}, \end{aligned} \tag{2}$$

where $cos(\mathbf{a},\mathbf{b})$ is the standard Cosine Similarity function. It can be seen that the relationship between $tri(\mathbf{a},\mathbf{b})$ and $cos(\mathbf{a},\mathbf{b})$ is simply a *consecutive* and *bijective* function[1] $f(z) = \sqrt{(1+z)/2}$ in its effective domain $z \geq -1$, indicating their equivalence in measuring a similarity.

A more intuitive interpretation of the relationship between the two similarities is illustrated in Fig. 1. In a Cartesian coordinate system, $\frac{\mathbf{a}}{\|\mathbf{a}\|}$ and $\frac{\mathbf{b}}{\|\mathbf{b}\|}$ represent two vectors lying on the unit circle, i.e. a circle of radius one. Their sum determines a directed chord (the blue line) on the circle, denoted by a new vector $\hat{\mathbf{c}}$. The Triangular Similarity is simply the half length of the chord, and the Cosine Similarity calculates the *cosine* of the angle θ between the two vectors. Note that the three vectors $\frac{\mathbf{a}}{\|\mathbf{a}\|}$, $\frac{\mathbf{b}}{\|\mathbf{b}\|}$ and their sum $\hat{\mathbf{c}}$ compose an *isosceles triangle*, that is why we call this similarity measurement as the *Triangular Similarity*.

2.2 Triangular Loss Function

For both the Cosine Similarity and the Triangular Similarity, the scale of $\|\mathbf{a}\|$ and $\|\mathbf{b}\|$, i.e. the length of the vectors, should be taken care of. Otherwise, it may raise up numerical instability problems [20] especially when $\|\mathbf{a}\|$ and $\|\mathbf{b}\|$ are too small. This hidden problem can be avoided or relieved by many strategies. For example, adopting regularization terms to prevent $\|\mathbf{a}\|$ and $\|\mathbf{b}\|$ from degenerating to 0 [5], normalizing the inputs by a whitening transformation [21, 43] so that all the input variables have unit variance [27], or simply performing an L2 normalization to let all the vectors have unit length [7, 38]. Empirical experiences showed that length regularization enables faster convergence [27] and better performance [7, 43] to machine learning applications.

[1] In mathematics, a bijective function is a function between the elements of two sets, where every element of one set is paired with exactly one element of the other set, and every element of the other set is paired with exactly one element of the first set.

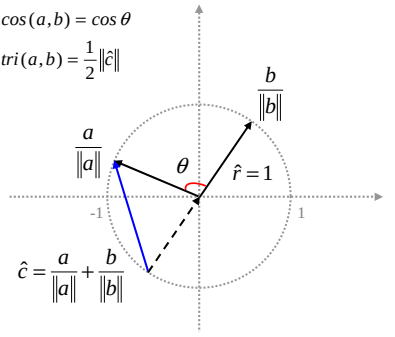


Figure 1: Equivalence between Triangular Similarity and Cosine Similarity. While the Cosine Similarity simply calculates the *cosine* of the angle θ between the two vectors, the Triangular Similarity measures the half length of the directed chord (the blue line). The three vectors $\frac{\mathbf{a}}{\|\mathbf{a}\|}$, $\frac{\mathbf{b}}{\|\mathbf{b}\|}$ and $\hat{\mathbf{c}}$ compose an *isosceles triangle*. The two similarity functions compose a one-to-one correspondence thus the equivalence is confirmed.

We propose a *soft L2 normalization* to constrain the length of the vectors. Instead of performing normalization as a preprocessing step on the inputs, we constrain the length of the vectors to a constant r by a regularization term:

$$\min \quad (\|\mathbf{a}\| - r)^2 + (\|\mathbf{b}\| - r)^2. \tag{3}$$

Minimizing the above function is able to make the values of $\|\mathbf{a}\|$ and $\|\mathbf{b}\|$ approaching a predefined scalar r. But in fact not all the vector lengths can be exactly r, so it is regarded as a *soft* length normalization.

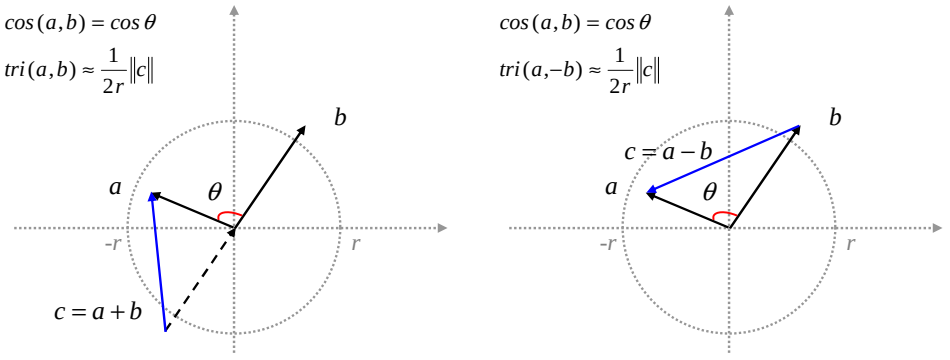
With the indicative assumption of $\|\mathbf{a}\| \approx \|\mathbf{b}\| \approx r$, both the Cosine Similarity and the Triangular Similarity can be simplified. For the Cosine Similarity, it can be rewritten as: $cos(\mathbf{a},\mathbf{b}) = \frac{\mathbf{a}^T\mathbf{b}}{\|\mathbf{a}\|\|\mathbf{b}\|} \approx (\frac{1}{r})^2\mathbf{a}^T\mathbf{b}$, where $\mathbf{a}^T\mathbf{b}$ is the *bilinear similarity*. This approximation explains that when vectors have approximate lengths, the bilinear similarity is equally effective but simpler than the Cosine Similarity in practical applications [7, 9, 13].

Similarly, we obtain a simplified version of Triangular Similarity as $tri(\mathbf{a},\mathbf{b}) = \frac{1}{2}\|\frac{\mathbf{a}}{\|\mathbf{a}\|} + \frac{\mathbf{b}}{\|\mathbf{b}\|}\| \approx \frac{1}{2r}\|\mathbf{a}+\mathbf{b}\|$. Instead of the isosceles triangle in Fig. 1, this simplified Triangular Similarity concerns a normal triangle lying around a circle of radius r (see Fig. 2 (a)). This triangle is determined by the two vectors $\mathbf{a}$ and $\mathbf{b}$, completed with their sum $\mathbf{c} = \mathbf{a}+\mathbf{b}$ as the third side. Additionally, the two vectors $\mathbf{a}$ and $\mathbf{b}$ determine another triangle where the third side is their difference, i.e. $\mathbf{c} = \mathbf{a}-\mathbf{b}$ (Fig. 2 (b)). With respect to the operations of sum or subtraction, we name the two triangles as the positive triangle and the negative triangle, respectively.

Like all the similarity metric learning methods, the objective of learning a Triangular Similarity metric is to increase the similarity between a similar pair and to decrease the similarity between a dissimilar pair. The geometrical interpretation of the objective is the following:

- When $\mathbf{a}$ and $\mathbf{b}$ are labeled as being similar, increasing the pairwise similarity means minimizing the inter-vector angle θ, which can be realized by maximizing the length of the third side $\mathbf{c}$ in the positive triangle (Fig. 2 (a)).

- When $\mathbf{a}$ and $\mathbf{b}$ are a dissimilar pair, we need to separte the two vectors with a larger angle θ. This can be also achieved by maximizing the length of the third side $\mathbf{c}$ in the negative triangle (Fig. 2 (b)).



(a) Positive Triangle (b) Negative Triangle

Figure 2: The simplified Triangular Similarity concerns normal triangles lying around a circle of radius r. A pair of vectors $\mathbf{a}$ and $\mathbf{b}$ determines two triangles: the positive triangle (left) illustrates the similarity between $\mathbf{a}$ and $\mathbf{b}$; the negative triangle illustrates the similarity between $\mathbf{a}$ and $-\mathbf{b}$.

The cost function

For a set of training data, let $s_i = 1$ (respectively -1) denote a pair of vectors $\mathbf{a}_i$ and $\mathbf{b}_i$ being similar (respectively dissimilar). With the soft length normalization factors $(\|\mathbf{a}_i\| - r)^2$, $(\|\mathbf{b}_i\| - r)^2$, and the lengths of the three sides, $\|\mathbf{a}_i\|$, $\|\mathbf{b}_i\|$, $\|\mathbf{c}_i\|$, the triangular loss of this pair is defined as:

$$J_i = \frac{1}{2}\Big[\boxed{(\|\mathbf{a}_i\| - r)^2 + (\|\mathbf{b}_i\| - r)^2}\Big] + r\Big(\boxed{\|\mathbf{a}_i\| + \|\mathbf{b}_i\| - \|\mathbf{c}_i\|}\Big), \tag{4}$$

where r is a constant constraint for the vector length; $\mathbf{c}_i = \mathbf{a}_i + s_i\mathbf{b}_i$, representing the simplified Triangular Similarity in a positive triangle or a negative triangle. It is interesting to find that the second part of this equation obeys the *triangle inequality theorem*. Moreover, the coefficients of the two parts are set to $\frac{1}{2}$ and r, in order to further simplify the formulation of Equation (4) as [45]:

$$J_i = \frac{1}{2}\|\mathbf{a}_i\|^2 + \frac{1}{2}\|\mathbf{b}_i\|^2 - r\|\mathbf{c}_i\| + r^2. \tag{5}$$

The gradient function

In Metric Learning systems, the vectors $\mathbf{a}_i$ and $\mathbf{b}_i$ are outputs of a certain mapping function $f(\cdot)$ parameterized by a set of parameters, e.g. $\mathbf{a}_i = f(\mathbf{x}_i, \mathbf{W})$ where $\mathbf{x}_i$ is an input vector. With respect to the parameter set $\mathbf{W}$, the derivative of a vector norm is known as $\frac{\partial\|\mathbf{a}\|}{\partial\mathbf{W}} = \frac{\partial\|\mathbf{a}\|}{\partial\mathbf{a}}\frac{\partial\mathbf{a}}{\partial\mathbf{W}} = \frac{\mathbf{a}}{\|\mathbf{a}\|}\frac{\partial\mathbf{a}}{\partial\mathbf{W}}$. Thus the derivative of the triangular loss (Equation (5)) is:

$$\frac{\partial J_i}{\partial\mathbf{W}} = (\mathbf{a}_i - r\frac{\mathbf{c}_i}{\|\mathbf{c}_i\|})\frac{\partial\mathbf{a}_i}{\partial\mathbf{W}} + (\mathbf{b}_i - r\frac{s_i\mathbf{c}_i}{\|\mathbf{c}_i\|})\frac{\partial\mathbf{b}_i}{\partial\mathbf{W}}. \tag{6}$$

The partial derivatives $\frac{\partial\mathbf{a}_i}{\partial\mathbf{W}}$ and $\frac{\partial\mathbf{b}_i}{\partial\mathbf{W}}$ are determined by the specific mapping function $f(\cdot)$, so the ideal minimal cost can be obtained at the zero gradient when $\mathbf{a}_i = r\frac{\mathbf{c}_i}{\|\mathbf{c}_i\|}$ and $\mathbf{b}_i = r\frac{s_i\mathbf{c}_i}{\|\mathbf{c}_i\|}$. In other words, the gradient function has $r\frac{\mathbf{c}_i}{\|\mathbf{c}_i\|}$ and $r\frac{s_i\mathbf{c}_i}{\|\mathbf{c}_i\|}$ as targets for $\mathbf{a}_i$ and $\mathbf{b}_i$, respectively.

The gradient function supports our predefined objective of learning the similarity, i.e., although the Triangular Similarity in the cost function is only a simplified version as we

have assumed the vectors having approximate lengths, we can still achieve that: (1) for two similar vectors, the gradient defines an identical target between them, i.e. $\mathbf{a}_i = \mathbf{b}_i = r\frac{\mathbf{c}_i}{\|\mathbf{c}_i\|}$; (2) for two dissimilar vectors, the gradient projects them to opposite directions, i.e. $\mathbf{a}_i = r\frac{\mathbf{c}_i}{\|\mathbf{c}_i\|}$ and $\mathbf{b}_i = -r\frac{\mathbf{c}_i}{\|\mathbf{c}_i\|}$.

2.3 Relation to Traditional Neural Networks

For classification problems, it is popular to use the Mean Squared Error (MSE) cost function in traditional neural networks, either an MLP [36] or a CNN [25]. It simply measures the difference between a computed output of a network and its desired target. Formally, for a given training sample $\mathbf{x}_i$ and its predefined target $\mathbf{g}_i$, we first compute its output by the mapping function, i.e. $\mathbf{a}_i = f(\mathbf{x}_i, \mathbf{W})$. The error for this training sample is the squared Euclidean distance between $\mathbf{a}_i$ and $\mathbf{g}_i$. The cost function and its gradient are:

$$J_i = \frac{1}{2}(\mathbf{a}_i - \mathbf{g}_i)^2; \quad \frac{\partial J_i}{\partial \mathbf{W}} = (\mathbf{a}_i - \mathbf{g}_i)\frac{\partial \mathbf{a}_i}{\partial \mathbf{W}}. \tag{7}$$

One can see that the gradient function of the triangular loss (Equation (6)) is exactly *a double copy* of the MSE gradient. The difference is that (1) the single output $\mathbf{a}_i$ in traditional neural networks is now paired with a partner $\mathbf{b}_i$ to learn the pairwise relationship between data; (2) the handcrafted target $\mathbf{g}_i$ is replaced by temporal targets $r\frac{\mathbf{c}_i}{\|\mathbf{c}_i\|}$ and $r\frac{s_i\mathbf{c}_i}{\|\mathbf{c}_i\|}$ which are automatically specified by the two vectors $\mathbf{a}_i$ and $\mathbf{b}_i$ themselves. This is indeed an advantage that the dimension of the output vectors is no longer required to be equal to the number of classes, and thus flexible dimensionality reduction is achieved by our metric learning system. Furthermore, with the similar gradient formulations, typical optimization techniques and practical tricks of training neural networks [33] – such as Weight Decay, Momentum [27], Dropout [39] – can be directly applied to optimize our system.

2.4 Comparison to Euclidean Distance

When we know that the Triangular Similarity is equivalent to the Cosine Similarity in measuring a data pair, one may care about the comparison between the Triangular Similarity and the Euclidean distance. The Euclidean distance is naturally related to the Cosine Similarity: for the Euclidean distance between two vectors, $\|\mathbf{a}-\mathbf{b}\|$, we know that $\|\mathbf{a}-\mathbf{b}\|^2 = (\mathbf{a}-\mathbf{b})^T(\mathbf{a}-\mathbf{b}) = \mathbf{a}^2 + \mathbf{b}^2 - 2\mathbf{a}^T\mathbf{b}$. If the vectors are normalized to unit length, i.e. $\mathbf{a}^2 = \mathbf{b}^2 = 1$, this equation reads that $\|\mathbf{a}-\mathbf{b}\|^2 = 2 - 2cos(\mathbf{a},\mathbf{b})$. That means, since we have controlled the vector length in our triangular loss function, minimizing the pairwise similarity value is approximately equivalent to maximizing the Euclidean distance.

However, whether the vector length is controlled or not, the advantage of the Cosine Similarity (as well as the Triangular Similarity) over the Euclidean distance is that it does not use the length information of vectors to distinguish two vectors. In other words, only the degree of freedom on the angle is enough to represent and distinguish vectors in different classes. In our similarity metric learning system, an operation of length normalization can be therefore adopted to removes length information of output vectors. This operation is especially interesting and useful for data visualization. Without the degree of freedom on the length, outputs are mapped to a sphere or hypersphere in the target space. According to the manifold learning theory that a hypersphere with a point removed is homeomorphic with a hyperplane [29], we can unfold the hypersphere into a lower dimensional space so that we obtain a new perspective to view the mapping result. This unfolding is simply realized by *coordinate transformation*, in a similar way of drawing a world map for the earth.

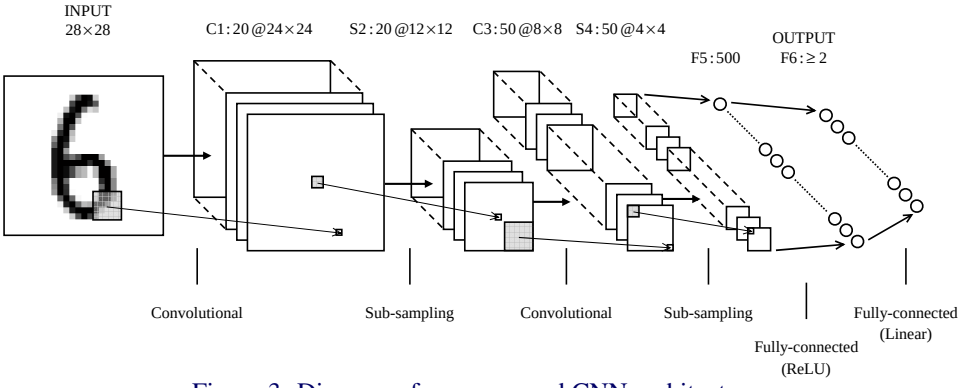


Figure 3: Diagram of our proposed CNN architecture.

3 End-to-end Data Visualization on Large-scale Data

3.1 The MNIST Dataset

The MNIST handwritten digits dataset [25][2] is a popular benchmark for classification by neural network based algorithms. There are 70,000 8-bit grayscale images in total, where 60,000 images are separated as the training data and the remaining 10,000 images are used for testing. All the images are of the same size 28×28, capturing a digit from 0 to 9 with various writing styles.

3.2 TSML-CNN

CNNs are a specialized kind of neural networks for processing data that have a known *grid-like topology* [9], i.e. 2-dimensional data such as images and speech time-series [24, 25, 26, 32]. In this work, we employ the CNN architecture recommended by [22][3] to process the images of the MNIST dataset and train our TSML system. In particular, this CNN architecture is similar with that in LeNet-5 [25] but of different hyperparameters. Figure 3 illustrates the layers in this CNN.

This CNN is composed of 6 layers, not counting the input layer. There are two convolutional layers at the 1st and 3rd layers, denoted by C1 and C3. Particularly, C1 and C3 are followed by two sub-sampling layers (i.e. pooling layers) S2 and S4. For all the convolutional layers, the size of local receptive fields is always 5×5. And a feature map is learned on one or more pages of units in the previous layer. For all the sub-sampling layers, the size of pools is always 2×2. Especially, we adopt a nonparametric sub-sampling operation, i.e. the max pooling: the four inputs in a pool are compared and the maximum is delivered to the following layer. The 5th layer F5 is a fully-connected layer having a ReLU activation function [15]. It contains 500 hidden units that are connected to every hidden unit in S4. The 6th layer F6 is also a fully-connected layer but performing a simple linear mapping into the target space. The size of this layer depends on the choice of the target space. For example, if we want to realize a mapping into the 3-dimensional space, we set the size of this layer to 3. We regard this layer as the output layer as we use the element values in this layer as the final vector representation for the input image.

[2] http://yann.lecun.com/exdb/mnist/
[3] http://caffe.berkeleyvision.org/

Let a matrix $\mathbf{X}$ denote an input image and a vector $\mathbf{a}$ represent the output vector, the above CNN accomplishes a deep nonlinear mapping $f(\cdot)$ that $\mathbf{a} = f(\mathbf{X}, \mathbf{W})$ where $\mathbf{W}$ indicates all the parameters in this CNN. Involving the CNN in our siamese architecture, TSML-CNN simply takes a pair of images $\mathbf{X}_i$ and $\mathbf{Y}_i$ from a training set and produces two outputs $\mathbf{a}_i = f(\mathbf{X}_i, \mathbf{W})$ and $\mathbf{b}_i = f(\mathbf{X}_i, \mathbf{W})$. TSML-CNN calculates the triangular loss on the two outputs by Equation (5). However, training a siamese network is much slower than training a traditional single-track network of the same size [44]. Especially, in the large training set of the MNIST dataset, the 60,000 training samples indicate 1.8 billion training pairs in total, thus the time consumption of directly training TSML-CNN by data pairs is unaffordable.

3.3 Hybrid Training: TSML-Hybrid

To provide an efficient dimensionality reduction, we propose a hybrid training algorithm, namely TSML-Hybrid that includes the following four stages:

1) *Tiny-scale training*: select only a few training samples from each class and train TSML-CNN on corresponding similar and dissimilar training pairs. In practice, the similar pairs play an important role in controlling vector lengths, thus we usually select at least 2 training samples from each class in order to generate both similar and dissimilar pairs. On the few training data, TSML-CNN gets convergence to an optimal solution very quickly.

2) *Transplanting*: set the centers of different classes in the target space as new labels for each class; take the learned parameters in TSML-CNN as initialization for a new CNN. This operation can be considered as transplanting the CNN in TSML-CNN to a new single-track CNN. With this transplantation, the single-track CNN needs no handcrafted target vectors and thus inherits the ability of flexible dimensional reduction from TSML-CNN.

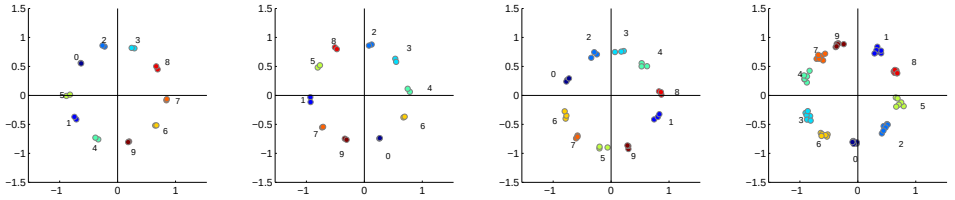
3) *Large-scale training*: take the new labels as target vectors for each class and train the single-track CNN on all training samples. The MSE function computes cost between output vectors and their desired targets. Compared with the large number of possible data pairs, the number of data samples is much smaller thus efficient convergence can be reached.

4) *Length normalization*: apply L2 normalization to make all the output vectors having unit length in the target space, because the Cosine Similarity (as well as the Triangular Similarity) does not need the length information of vectors to distinguish different classes.

Compared to directly training a siamese network such as TSML-CNN, the proposed hybrid algorithm greatly improves the efficiency of training but does not suppress the classification performance. We will provide more results in the experimental section.

4 Results and Analysis

In this section, we visualize the MNIST digits in the 2/3/4-dimensional spaces by using HSML-Hybrid.



(a) 20 samples (b) 20 samples (c) 30 samples (d) 60 samples

Figure 4: Results after tiny-scale training with different data size and initialization.

4.1 Results of Data Visualization

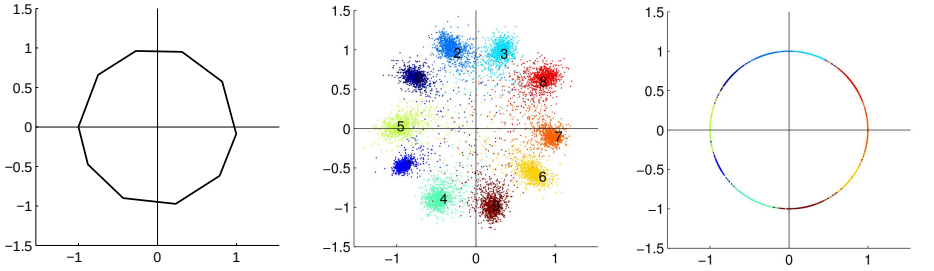
4.1.1 Case study 1: 2-dimension

In *tiny-scale training*, we randomly select 2 training sample from each of the 10 classes in the MNIST dataset. We train TSML-CNN on these data pairs with the output dimension equal to 2. The learned model projects the 20 training samples into the target space as in Fig. 4 (a), where numbers beside the points denote the digits they represent. The positions of the digits are determined by the random initialization and tiny-scale training procedure. We observe no evident relation between adjacent digits in our experiments: Fig. 4 shows other mapping results with different data size and initialization, though the order of the 10 digits varies, the overall distribution is always approximating a circle.

After the tiny-scale training, the 20 samples evenly distribute along a circle around the origin, connecting the centers of every class results in a 10-sided polygon (Fig. 5 (a)). In *large-scale training*, the class centers are set as labels (i.e. target vectors) for their corresponding classes. Minimizing the MSE cost function attracts all the training samples to these centers. The final distribution of the 10,000 test data after this training is shown in Fig. 5 (b). We can see that data of each class scatter in a local small area in the 2-dimensional space, assembling around the labeled class center. By length normalization, the data vectors are then projected to the unit circle (see Fig. 5 (c)).

In manifold learning theory, a circle with a point removed is homeomorphic with a real line. In other words, the points on the circle can be represented by the degree of freedom on the angle only. In Figure 6, we unfold the circle to an interval $[0, 2\pi]$ in the lower picture, that means each test image is represented by a single point within this interval. In the upper picture, the histograms of the test data in the 10 classes are shown. We can see that each class has a sharp peak and adjacent classes are divided with clear valleys, i.e. data in the same class have been well concentrated and data in different classes have been well separated. One may notice the shape of different histograms varies, e.g. the peak of the number 4 is low and fatty but that of the number 1 is high and sharp. We found that the shape of the histogram – the final distribution of the output data – is automatically determined by the random initialization and local optimal convergence of the present neural networks.

However, recent theoretical and empirical results strongly suggest that the different local optima are not a serious issue in general [28]. Regardless of initial conditions that generate different output distributions, our non-convex system nearly always reaches local solutions of very similar classification quality. Specifically, kNN classification on either the 2-dimensional data on the circle or the 1-dimensional data within the interval yields an accuracy of 94.76%. In conclusion, with only one mapping by the proposed similarity metric learning, we have obtained two views for visualizing the test data, i.e. Fig. 5 (c) and Fig. 6.



(a) (b) (c)

Figure 5: After tiny-scale training, (a) centers of different classes stand as vertexes of a polygon; after large-scale training, (c) the 10,000 test data assemble around each center; by length normalization, (d) the data are further projected onto the unit circle.

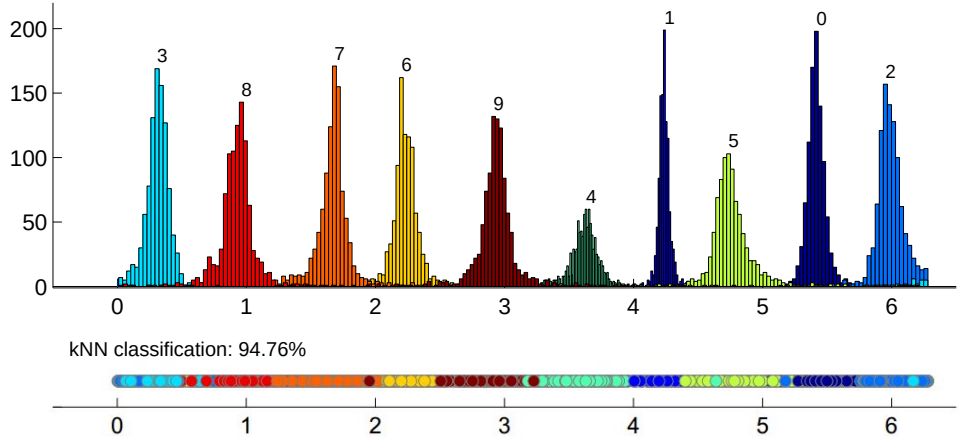
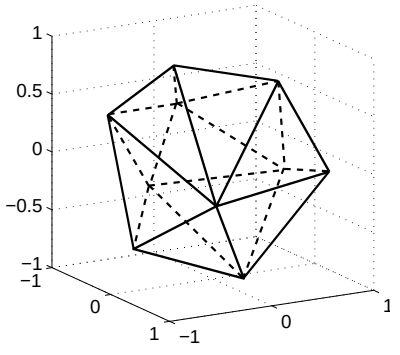


Figure 6: The lower picture gives an unfolded view of the MNIST test data in Fig. 5 (c), i.e. an interval $[0, 2\pi]$; the upper picture plots the histograms of the 10 classes.

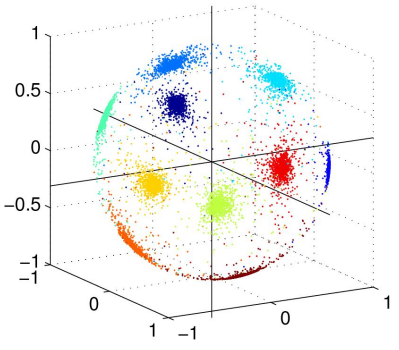
4.1.2 Case study 2: 3-dimension

When we set the output dimension to 3, a different model is learned on the same 20 training samples. Firstly, the well-trained TSML-CNN projects the 20 training samples into the 3-dimensional space, comprising a convex polyhedron with 10 vertexes that will be later used as target vectors for the 10 classes, as shown in Fig. 7 (a).

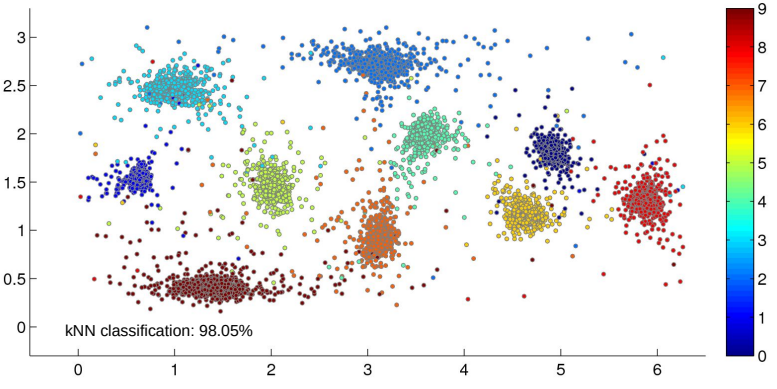
Similar with the unit circle in the case of 2-dimension, the TSML-Hybrid model produces a unit sphere carrying all the 3-dimensional data (see Fig. 7 (b)). More interestingly, a sphere with a point removed is homeomorphic with a plane. In other words, considering the sphere as the earth, the plane is like the world map. Figure 7 (c) illustrates the unfolded map for the test data. On this map, different classes of data locate as 10 isolated continents. Note that the leftmost part of the map and the rightmost part are actually connected in the 3-dimensional space. The length and width of this map are 2π and π, respectively. In addition, kNN classification accuracy of the 2-dimensional data is 98.05%, much higher than that of the 1-dimensional data (94.76%).



(a) Polyhedron connecting the target centers



(b) Test data after large-scale training



(c) Unfolded view of the test data

Figure 7: Mapping results in the 3-dimensional target space.

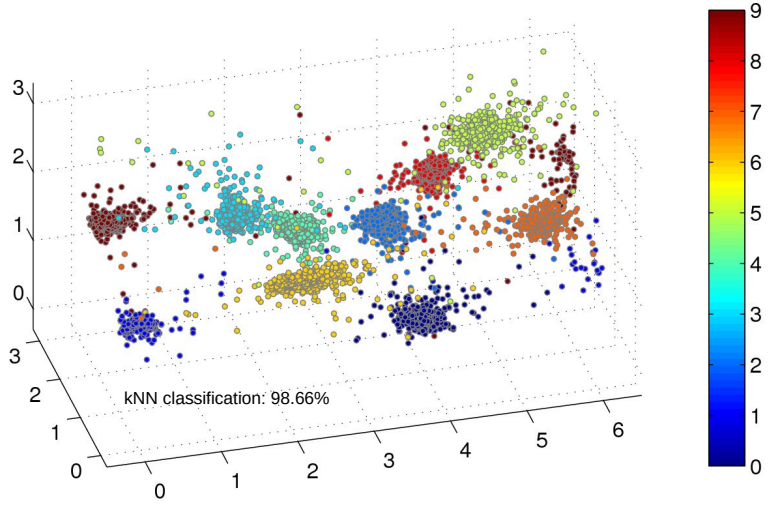
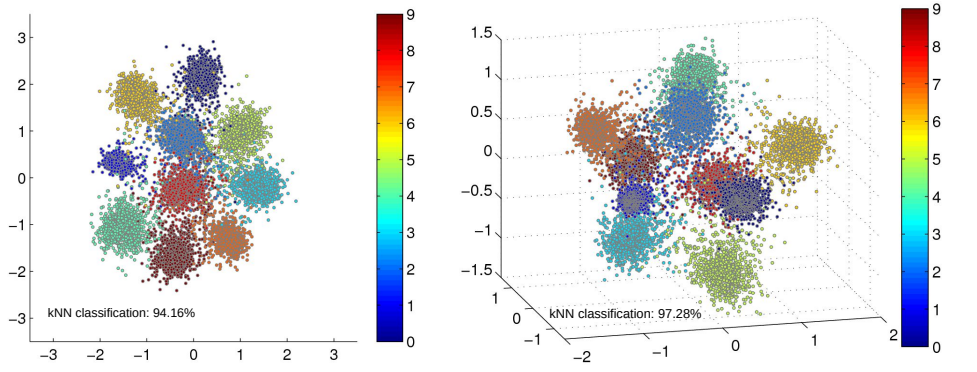


Figure 8: The unfolded view of the 4-dimensional MNIST test data in the 3-dimensional space.



(a) In 2-dimensional space (b) In 3-dimensional space

Figure 9: Mapping results by DrLIM in (a) 2-dimensional and (b) 3-dimensional spaces.

4.1.3 Case study 3: 4-dimension

In fact, we can not see the 4-dimensional space but only a subspace of it. Fortunately, the proposed method provides such a subspace of the target space, i.e. the unit hypersphere. Unfolding this 4-dimensional hypersphere results in a part of the 3-dimensional space (see Fig. 8). The length, width and height of this partial space are 2π, π and π, respectively. Like the previous cases, we observe clear boundaries between different classes. One may notice that a few data of the two classes at the leftmost side are mapped to the rightmost side, actually they are together in the 4-dimensional space. In terms of better data visualization, a probable solution can be duplicating the 3-dimensional mapping results to make the visualization locally complete. In terms of classification, since the *k*NN classifier makes prediction on the majority label of *k*-nearest neighbors in the training set, this phenomenon of separated data has little influence on the classification performance. The *k*NN classifier obtains an classification accuracy of 98.66% on the unfolded 3-dimensional data.

4.2 Comparison to the State-of-the-art

The most related work with ours is Dimensionality Reduction by Learning an Invariant Mapping (DrLIM) [19], which proposed a contrastive loss function to measure the cost of data pairs. Different with our triangular loss function concerning the Triangular Similarity (or the Cosine Similarity), the contrastive loss function was based on the Euclidean distance. The objective of the contrastive loss is to minimize the distance between a similar pair and to separate any two dissimilar data with a distance margin. With the same CNN architecture used in TSML-Hybrid, the implemented DrLIM in [22] projects the MNIST test data into the 2-dimensional and 3-dimensional spaces as in Fig. 9. DrLIM does not benefit from the manifold learning theory and thus provides only one perspective to view the outputs.

Running time: using the deep learning tool Caffe [22] in CPU mode only, we carried out all the experiments on a machine with a 2.3 GHz dual-core CPU, 4GB RAM and 64-bit operating system. In general, as DrLIM and TSML-Hybrid have the same deep CNN architecture, performing gradient descent once for either the triangular loss or the contrastive loss costs almost the same time. However, the hybrid training strategy helps TSML-Hybrid to train a model much faster. Concretely, the tiny-scale training costs about 150 seconds and the large-scale training costs about 960 seconds, hence obtaining the TSML-Hybrid models averagely costs about 1110 seconds (see Table 1). For DrLIM that directly performs large-

Table 1: Comparison on training time (in seconds) of different methods. '—' means no available result. View 1: before unfolding; View 2: after unfolding.

Methods	1-dimension	2-dimension	3-dimension	10-dimension	Mean
CNN classifier [22]	—	—	—	968.33	968.33
DrLIM [19]	—	9632.10	9579.83	9556.59	9589.51
TSML-Hybrid (view 1)	—	1111.64	1113.89	1119.19	1116.75
TSML-Hybrid (view 2)	1111.64	1113.89	1121.10	1117.92	

Table 2: Comparison on classification accuracy of different methods. '—' means no available result. View 1: before unfolding; View 2: after unfolding.

Methods	1-dimension	2-dimension	3-dimension	10-dimension
CNN classifier [22]	—	—	—	98.97%
DrLIM [19]	—	94.16%	97.28%	98.35%
TSML-Hybrid (view 1)	—	94.76%	98.05%	**99.02%**
TSML-Hybrid (view 2)	**94.76%**	**98.05%**	**98.66%**	98.87%

scale training on data pairs, the convergence speed slows down significantly than the other methods, the average training time for each DrLIM model is about 9600 seconds.

Classification accuracy: a comparison on classification performance is summarized in Table 2. The classical CNN classifier – the present CNN followed by a Euclidean loss layer or softmax loss layer – is set as the baseline for classification. The proposed TSML-Hybrid method offers two views of the target space. Take a target space of dimension 10 for example, View 1 is the result of TSML-Hybrid with the output layer size equal to 10; View 2 is the result of an output dimension 11 and then unfolded into the 10-dimensional space by coordinate transformation.

Firstly, we evaluate all the methods in the 10-dimensional space because the output of CNN is determined by the number of classes. Comparing these results we find that our TSML-Hybrid method obtains comparable classification results with the CNN classifier (98.97%) in the 10-dimensional space (99.02% and 98.87% in the two views respectively). However, DrLIM relatively performs worse (98.35%) than the other three. This is because DrLIM trains on a limited number of data pairs and does not fully make use of all the labeled data samples.

Concerning the lower dimensional spaces, the CNN classifier failed for dimensionality reduction because the dimension of the handcrafted target vectors was fixed. DrLIM and TSML-Hybrid relaxes this constraint by metric learning and realizes classification in the visualizable spaces. Generally, TSML-Hybrid (View 2) gets better results than DrLIM and TSML-Hybrid (View 1) in these low dimensional spaces. For DrLIM and TSML-Hybrid (View 1), the classification accuracy decreases as the target dimension is reduced and a large decline occurs when the target dimension changes from 3 to 2. By using coordinate transformation, TSML-Hybrid (View 2) delays the decline till the dimension reduced from 2 to 1. Therefore, the View 2 is preferred if the classification performance is mainly concerned.

5 Conclusion

In this paper, we introduced a novel Triangular Similarity to define the pairwise data relationship in the target spaces and deploy it into the CNN architecture to project an image to a point in low dimensional spaces. Interestingly, this similarity metric learning approach

enables us to make use of manifold learning theories and visualize the output data in two different views. We succeeded in mapping thousands of images to local parts of low dimensional spaces but having a competitive classification accuracy. In the near future, we would like to apply the method on other datasets and extend our work to unsupervised tasks.

References

[1] Pierre Baldi and Yves Chauvin. Neural networks for fingerprint recognition. *Neural Computation*, 5(3):402–418, 1993.

[2] Aharon Bar-Hillel, Tomer Hertz, Noam Shental, and Daphna Weinshall. Learning a mahalanobis metric from equivalence constraints. *Journal of Machine Learning Research*, 6(6):937–965, 2005.

[3] Aurélien Bellet, Amaury Habrard, and Marc Sebban. A survey on metric learning for feature vectors and structured data. *arXiv preprint arXiv:1306.6709*, 2013.

[4] Yoshua Bengio, Ian J. Goodfellow, and Aaron Courville. Deep learning. Book in preparation for MIT Press, 2015. URL http://www.iro.umontreal.ca/~bengioy/dlbook.

[5] Samuel Berlemont, Gregoire Lefebvre, Stefan Duffner, and Christophe Garcia. Siamese neural network based similarity metric for inertial gesture classification and rejection. In *11th IEEE International Conference on Automatic Face and Gesture Recognition*, May 2015.

[6] Jane Bromley, James W Bentz, Léon Bottou, Isabelle Guyon, Yann LeCun, Cliff Moore, Eduard Säckinger, and Roopak Shah. Signature verification using a "siamese" time delay neural network. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(04):669–688, 1993.

[7] Qiong Cao, Yiming Ying, and Peng Li. Similarity metric learning for face recognition. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 2408–2415. IEEE, 2013.

[8] Gal Chechik, Uri Shalit, Varun Sharma, and Samy Bengio. An online algorithm for large scale image similarity learning. In *Advances in Neural Information Processing Systems*, pages 306–314, 2009.

[9] Gal Chechik, Varun Sharma, Uri Shalit, and Samy Bengio. Large scale online learning of image similarity through ranking. *The Journal of Machine Learning Research*, 11: 1109–1135, 2010.

[10] Ke Chen and Ahmad Salman. Extracting speaker-specific information with a regularized siamese deep network. In *Advances in Neural Information Processing Systems*, pages 298–306, 2011.

[11] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *Computer Vision and Pattern Recognition, IEEE Conference on*, volume 1, pages 539–546. IEEE, 2005.

[12] Jason V Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S Dhillon. Information-theoretic metric learning. In *Proceedings of the 24th International Conference on Machine Learning*, pages 209–216. ACM, 2007.

[13] Jia Deng, Alexander C Berg, and Li Fei-Fei. Hierarchical semantic indexing for large scale image retrieval. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 785–792. IEEE, 2011.

[14] Amir Globerson and Sam T Roweis. Metric learning by collapsing classes. In *Advances in neural information processing systems*, pages 451–458, 2005.

[15] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *International Conference on Artificial Intelligence and Statistics*, pages 315–323, 2011.

[16] Jacob Goldberger, Geoffrey E Hinton, Sam T Roweis, and Ruslan Salakhutdinov. Neighbourhood components analysis. In *Advances in neural information processing systems*, pages 513–520, 2004.

[17] Alex Graves and Navdeep Jaitly. Towards end-to-end speech recognition with recurrent neural networks. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1764–1772, 2014.

[18] Matthieu Guillaumin, Jakob Verbeek, and Cordelia Schmid. Is that you? metric learning approaches for face identification. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 498–505. IEEE, 2009.

[19] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *Computer vision and pattern recognition, 2006 IEEE computer society conference on*, volume 2, pages 1735–1742. IEEE, 2006.

[20] Nicholas J Higham. *Accuracy and stability of numerical algorithms*. Siam, 2002.

[21] Aapo Hyvärinen and Erkki Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4):411–430, 2000.

[22] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*, pages 675–678. ACM, 2014.

[23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[24] Steve Lawrence, C Lee Giles, Ah Chung Tsoi, and Andrew D Back. Face recognition: A convolutional neural-network approach. *Neural Networks, IEEE Transactions on*, 8 (1):98–113, 1997.

[25] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[26] Yann LeCun, Koray Kavukcuoglu, and Clément Farabet. Convolutional networks and applications in vision. In *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*, pages 253–256. IEEE, 2010.

[27] Yann A LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012.

[28] Yann A LeCun, Yoshua Bengio, and Geoffrey E Hinton. Deep learning. *Nature*, 521: 436–444, 2015.

[29] Bert Mendelson. *Introduction to topology*. Courier Corporation, 1990.

[30] Hossein Mobahi, Ronan Collobert, and Jason Weston. Deep learning from temporal coherence in video. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 737–744. ACM, 2009.

[31] Hieu V Nguyen and Li Bai. Cosine similarity metric learning for face verification. In *Computer Vision–ACCV 2010*, pages 709–720. Springer, 2011.

[32] Feng Ning, Damien Delhomme, Yann LeCun, Fabio Piano, Léon Bottou, and Paolo Emilio Barbano. Toward automatic phenotyping of developing embryos from videos. *Image Processing, IEEE Transactions on*, 14(9):1360–1371, 2005.

[33] Genevieve B Orr and Klaus-Robert Müller. *Neural networks: tricks of the trade*. Springer, 2003.

[34] Ali Mustafa Qamar and Eric Gaussier. Online and batch learning of generalized cosine similarities. In *Data Mining, 2009 IEEE International Conference on*, pages 926–931. IEEE, 2009.

[35] Ali Mustafa Qamar, Eric Gaussier, Jean-Pierre Chevallet, and Joo Hwee Lim. Similarity learning for nearest neighbor classification. In *Data Mining, 2008 IEEE International Conference on*, pages 983–988. IEEE, 2008.

[36] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, DTIC Document, 1985.

[37] Shai Shalev-Shwartz, Yoram Singer, and Andrew Y Ng. Online and batch learning of pseudo-metrics. In *Proceedings of the twenty-first international conference on Machine learning*, page 94. ACM, 2004.

[38] K. Simonyan, O. M. Parkhi, A. Vedaldi, and A. Zisserman. Fisher vector faces in the wild. In *British Machine Vision Conference*, 2013.

[39] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[40] Kilian Q Weinberger and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. *The Journal of Machine Learning Research*, 10:207–244, 2009.

[41] Eric P Xing, Michael I Jordan, Stuart Russell, and Andrew Y Ng. Distance metric learning with application to clustering with side-information. In *Advances in neural information processing systems*, pages 505–512, 2002.

[42] Wen-tau Yih, Kristina Toutanova, John C Platt, and Christopher Meek. Learning discriminative projections for text similarity measures. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 247–256. Association for Computational Linguistics, 2011.

[43] Lilei Zheng, Cheung-Chi Leung, Lei Xie, Bin Ma, and Haizhou Li. Acoustic texttiling for story segmentation of spoken documents. In *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pages 5121–5124. IEEE, 2012.

[44] Lilei Zheng, Stefan Duffner, Khalid Idrissi, Christophe Garcia, and Atilla Baskurt. Siamese multi-layer perceptrons for dimensionality reduction and face identification. *Multimedia Tools and Applications*, 2015.

[45] Lilei Zheng, Khalid Idrissi, Christophe Garcia, Stefan Duffner, and Atilla Baskurt. Triangular similarity metric learning for face verification. In *11th IEEE International Conference on Automatic Face and Gesture Recognition*, May 2015.